A Closed Form Solution for the Brightness Preserving Colour to Greyscale Image Conversion

Graham D. Finlayson and Toby N. Matheson; School of Computing Sciences; University of East Anglia; Norwich, UK

Abstract

There are many methods for converting a colour image to a grey scale counterpart. The luminance image can be calculated as a weighted sum of R, G and B. However, when equiluminant edges appear in images, they disappear in the greyscale reproduction. Alternate greyscale computations attempt to mitigate this problem by finding the best solution according to an optimisation criterion. Optimisations include best representing the colour difference in grey scale or maximising the variance of the greyscale reproduction. A promising previous approach proposed maximising the contrast of a greyscale reproduction subject to the constraint that the brightness was preserved (i.e. the grey scale reproduction would have the same brightness as the colour original). The required greyscale was found using a quadratic programming optimisation. While this made the algorithm simple to describe it limited its practical utility (e.g. it is unlikely to get QP implemented in a digital camera). The main result of this paper is to show that there exists a closed form solution for finding the maximum contrast and brightness preserving greyscale.

As in the previous work, we define that a greyscale is a weighted sum of R, G and B, and that resulting greyscale has the same average as the colour original. We propose that the individual weights should be between 0 and 1 and their sum is equal to 1 (this constraint appeals to our notion of reasonableness and ensures white is preserved). These constraints coupled with our requirement that brightness is preserved is interpreted geometrically. We show that the vector of 3 weighting factors must lie on a line segment and that the best solution is always at one of the endpoints. It is straightforward to directly solve for these endpoints and so directly solve the maximum contrast brightness preserving greyscale problem.

Introduction

The simplest method for colour to greyscale conversion is the simple weighted summation:

$$L(x,y) = \alpha R(x,y) + \beta G(x,y) + \gamma B(x,y) \tag{1}$$

where R, G and B represent the red, green and blue channels and α, β and γ are their weightings respectively such that:

$$0 \leq \alpha \leq 1 \qquad 0 \leq \beta \leq 1 \qquad 0 \leq \gamma \leq 1 \tag{2}$$

subject to the constraint:

$$\alpha + \beta + \gamma = 1 \tag{3}$$

Here we assume that the RGB image is linear.

The constraints in (3) provide a reasonable definition of greyscale (a weighted sum where the sum of the weights is 1). One of the properties this greyscale transform defines, is that (1,1,1) white maps to 1 on the output (and this will have maximum brightness).

If we are in, for example, linear sRGB space [2] there is a defined α=0.2126, β=0.7152, γ=0.0722 which maps the RGB's to luminance. The luminance image is clearly a reasonable choice for colour to greyscale conversion. However, greyscale conversion based on luminance values at each pixel can be prone to a loss of information where pixels with differing chromaticities but identical luminances will appear – post conversion – as the same shade of grey. This observation, in part, has been the genesis of a large body of work that attempts to find a better way for converting colour into greyscale. 'Better' might mean that the colour differences in the original image are reflected in the greyscale reproduction (e.g. Gooch [7]). Alternatively, Bala [3] simply adds a chroma-type edge mask to the luminance image (ensuring that equiluminant edges do not disappear). However, there are several algorithms reported that attempt to maximise contrast in the greyscale reproduction (including Alsam [1] and Qiu et al. [10])).

Significantly, Connah et al. [4] investigated how a simple luminance encoding compared, from a preference standpoint, to other colour to greyscale algorithms. Given a set of algorithms – comprising Alsam [1], Grundland [8], Socolinsky [12], Rasche [11] and Bala [3] – Connah et al. found that one of these algorithms was *always* preferred to luminance.

In this paper, we build on the work by Qiu et al. [10]. That work addresses the contrast question head on, and presents a computational algorithm that delivers a greyscale image which for a fixed overall brightness has maximum contrast. They formulate the greyscale conversion problem as a quadratic program. The objective is to maximise the global contrast (a quadratic concept) subject to the linear constraint that global brightness is preserved. Unfortunately, quadratic programming (QP), though indubitably an extremely powerful optimisation method, is a complex algorithm. The main technical contribution of this paper is to show how we can solve the maximum contrast brightness preserving greyscale using simple analytical calculations.

In our approach we suppose that the created greyscale must be reasonable (it must follow Eq. (1)). That is the sum of the coefficients is 1 and each coefficient is larger than 0 and less than 1. We seek a greyscale where brightness is preserved which implies that the mean of the generated greyscale is the same as the mean of the original colour image. This is a second linear constraint on the coefficients we seek:

$$\alpha\mu_R + \beta\mu_G + \gamma\mu_B = \mu_g, \tag{4}$$

where μ_R, μ_G and μ_B are the respective means of the R, G and

B image channels. The right-hand side of the equation μ_g (is the mean of the original colour images i.e. sum up all of the N×M×3 pixel values and divide by 3NM).

Mathematically, each linear equation (3) and (4) restricts the vector of coefficients $[\alpha \quad \beta \quad \gamma]$ to lie on a plane. Since we have two linear constraints we have two planes and the coefficient vector must lie on the intersection of the two (e.g. a line). The greater than 0 and less than 1 constraints further restricts the solution set to lie on a line segment. We show that the maximum contrast solution is at one of the endpoints of this segment. Simple algebraic methods can be used to find the line segment and its endpoints. So, it is a simple matter to find the maximum contrast greyscale that preserves brightness.

By definition our method provides the same images as calculated by Qiu et al.'s QP algorithm [10]. We calculate greyscale outputs for some of the standard images used in the literature to evaluate the effectiveness of our algorithm. These reproductions are compared with the outputs of competing methods. Overall, our images have more contrast than the antecedent methods.

Background

It is immediate that in generating a greyscale image we are, fundamentally, reducing a 3-dimensional colour image to a 1-dimensional counterpart. Principal component analysis is the natural way of achieving such dimension reduction. For example, Alsam proposed to use PCA as a tool for generating greyscales [1]. The (non mean subtracted) covariance matrix (properly called the Raw Cross Product Matrix) of the colour image is calculated. The direction in colour space which has maximum variance (the first principal component direction) is found, and projecting the colours onto this axis results in a base greyscale image.

Alsam et al. [1] acknowledged that maximising variance alone did not guarantee that colour edges in the original image did not disappear in their reproduction. Thus an edge map computed from the colour original is superimposed on their base image.

In separate work, Bala et al. [3] also calculates a colour edge map and superimposes this onto the standard luminance image. (In Bala et al.'s work, effectively, the base image is always the luminance image).

Grundland and Dodgson [8] use a novel projection-based dimension reduction approach which they call 'Predominant Component Analysis'. The image is converted to YIQ space where the luminance space Y, and the chrominance space IQ can be separated and treated differently. They use the calculated 'chromatic and luminance differences' to define a predominant chromatic axis, which is then used as a basis onto which the chromatic differences are projected. This essentially transforms the IQ space into a single chromatic channel. This channel is then weighted and added onto the luminance channel thus producing a greyscale image.

Rasche, Geist and Westall [11] developed a contrast preserving algorithm that maintains luminance consistency expressed in terms of a constrained, multidimensional scaling (MDS) problem. Effectively, their method exploits the idea that a distance matrix can be computed for every colour in an image to every other colour. MDS is a means of finding a 1-dimensional representation such that the 3-d distances are preserved in the reduced space .

An elegant, local, contrast encoding algorithm for colour to greyscale was developed by Socolinsky and Wolff [12]. Their key insight is to show that the 3 colour derivatives in an image can be approximated by greyscale equivalent derivatives. That is the three x and y derivatives (one per channel) can be replaced by single approximate x and y derivatives (for the interested reader this is done using Di Zenzo's colour tensor decomposition [5]). The Socolinsky and Wolff grey scale is found by reintegrating these equivalent derivatives. Since the calculated derivatives may not actually correspond to any real image, the optimal solution is found by solving a Poisson's equation (i.e. choosing the image whose derivatives are closest to the ones we seek).

The recovered images in a strong mathematical sense are similar to the colour original. A disadvantage of Socolinsky and Wolff's approach [12] is that the recovered images can have gradients which were not in the original (the recovered images often suffer from bending and smearing artefacts).

Gooch and coworkers [7] have developed a local colour to greyscale that bears resemblance to Grundland and Dodgson's work [8], in that they split up their luminance and chrominance image information after converting it to a colour space (CIE L*a*b* in this case). The luminance and chrominance are used to set a target difference, which is then used in defining a quadratic objective function for finding their brightness image which they solve using standard techniques.

To some extent all the algorithms reviewed above change and generally enhance the contrast of the greyscale reproduction (compared to luminance). Yet, Qiu et al. [10] observed that these algorithms also change the image brightness. An image I that is multiplied by a scalar k (and assuming the resulting image has no clipped values) will have higher brightness but also higher contrast (at least as it is generally defined with respect to colour to greyscale computation). This led Qiu et al. [10] to novelly consider maximising contrast while maintaining the reproduction brightness.

The problem has a natural QP formulation since contrast is defined as image variance (squared or quadratic), and the idea that brightness is preserved is a linear constraint. Significantly, the added constraint of preserving brightness can result in dramatically different images than when this constraint is not enforced, and, informally at least, the reproductions have more image contrast. Unfortunately, QP is a complex computer algorithm, and so, although this algorithm provides compelling reproductions, it is not easily implementable (e.g. in an embedded device).

In Qiu et al.'s approach [10] contrast is defined to be the variance of the greyscale reproduction which can be computed directly from the covariance matrix Σ of the colour original:

$$\Sigma_{i,j} = E[(Q_i - \mu_i)(Q_j - \mu_j)] \qquad i,j \in \{R,G,B\} \tag{5}$$

Here E is the expectation operator, Q_i is a pixel value in a given channel and μ_i is the mean of that channel.

To ease the mathematical formulation we represent the weighting coefficients as a vector and define the auxiliary vectors u and μ:

$$\underline{v}^t = [\alpha \quad \beta \quad \gamma],\ \underline{u}^t = [1 \quad 1 \quad 1],\ \underline{\mu}^t = [\mu_R \quad \mu_G \quad \mu_B] \tag{6}$$

where the superscript t denotes a vector transpose.

We can now write the contrast maximising and brightness preserving greyscale problem as a quadratic optimisation with linear constraints:

$$\max_{\underline{v}} ||\underline{v}^t \Sigma \underline{v}|| \quad s.t. \quad \underline{v}^t \cdot \underline{u} = 1 \quad \& \quad \underline{v}^t \cdot \underline{\mu} = \mu_g \tag{7}$$

In Qu et al.'s original paper [10], this problem is solved using quadratic programming which is a general search based algorithm guaranteed to find the global optimum. A full description of QP can be found in any non-linear programming text [9].

A Closed Form Solution

Given the quadratic programming formulation in (7) we will now show how we can solve for the greyscale that maximises contrast and at the same time preserves brightness (see QP in Eq. (7)). Our aim is to do this without using the heavy and computationally cumbersome tool of quadratic programming. We formulate the problem in such a way that a natural closed form solution presents itself.

We illustrate our argument somewhat informally by appealing to the inherent geometry of the problem. A detailed and formal proof is presented in a companion paper. However, the inherent geometry is really quite simple and so by drawing a few well chosen figures we can effectively derive our result.

First we return to Eq. (2). Here we capture the notion that in forming a linear sum of R, G and B to form a greyscale that the coefficients are bigger than 0 and less than 1. Thus the coefficient vector (which we denoted $\underline{v}$ in (6)) must lie inside the unit cube, see Figure 1(a).

The idea of reasonable linear combination is defined in (3) i.e. the idea that the total of the weights must sum up to 1. If this were not the case, a white response triplet in the greyscale would be mapped darker or brighter than white. In the former case the resulting image would appear lacking in contrast and dull. In the latter we would have greyscale values we could not display. Eq. (3) in fact is the equation of a plane in coefficient space. We can define a plane given knowledge of 3 points that lie on the plane. Setting two of the three coefficient vectors to zero we find that (1,0,0), (0,1,0) and (0,0,1) must all lie on the plane. At the same time the part of plane we are interested in must lie inside the unit cube. This gives us the triangular region shown in Figure 1(b). Interestingly, by appealing to the reasonableness of the linear combination forming a greyscale, we already greatly restrict where our coefficient vector must lie.

The global mean i.e. the sum of all N, red pixels + N green pixels + N blue pixels (divided by 3N) is our definition of the brightness that must be preserved. Eq. (4) relates the global mean to the means in each of the red, green and blue channels. Eq. (4) is again an equation of a plane where the coefficient vector that we seek must lie. We illustrate a possible plane in Figure 1(c) (the orientation of which is defined by the relative means of the original colour image). Again we have a triangular region because we need to enforce the constraints that the coefficients lie between 0 and 1.

That our greyscale image is reasonable and that the global means are preserved implies that the coefficient vector we seek must, simultaneously, lie in both planes. The intersection of two planes is a line. For our example the plane intersection is shown in Figure 1(d). It is clear that the dotted line segment delimits all possible solutions to our brightness preserving grey scale problem.

Suppose we denote the coefficient vectors for the endpoints of this line as $\underline{x}$ and $\underline{y}$. Then the contrast for any point on the line between these points can be written as:

$$contrast = \underline{v}^t \Sigma \underline{v} \quad s.t. \quad \underline{v} = \rho \underline{x} + (1-\rho)\underline{y}, \quad \rho \in [0,1] \tag{8}$$

$$contrast = \rho^2 \underline{x}^t \Sigma \underline{x} + 2\rho(1-\rho)\underline{x}^t \Sigma \underline{y} + (1-\rho)^2 \underline{y}^t \Sigma \underline{y} \tag{9}$$

Since the vector $\underline{x}$ and $\underline{y}$ are known, it turns out that our measure of contrast is simply a quadratic equation in the mixing coefficient ρ. Like all quadratic equations there will be a single maximum or minimum stationary point. If it is a minimum then the solution we seek must be at the end points of the line segment, as it is at one of these points where (9) reaches a maximum value. In fact it is easy to show for the problem at hand that the coefficient vector does indeed lie at either end of the line segment,

$$A = \underline{x}^t \Sigma \underline{x} \qquad B = \underline{x}^t \Sigma \underline{y} \qquad C = \underline{y}^t \Sigma \underline{y} \tag{10}$$

$$contrast = (A - 2B + C)\rho^2 + 2(B - C)\rho + C \tag{11}$$

because this is a second order polynomial, we know (save degenerate cases) there is a single maximum or minimum. From calculus we know that if the 2nd derivative is larger than zero then the corresponding function has a local minimum. We differentiate (11) twice we find:

$$\frac{d^2 contrast}{d\rho^2} = 2(A - 2B + C) \tag{12}$$

Because the covariance matrix is positive semi definite (and α, β and γ are between 0 and 1) then A, B and C are all larger than or equal to zero.

Let us suppose that the covariance matrix Σ is the identity matrix. Then equation (12) is equal to,

$$2(|\underline{x}|^2 + |\underline{y}|^2 - 2|\underline{x}||\underline{y}|) \tag{13}$$

which is precisely the squared magnitude of the vector $\underline{x} - \underline{y}$. Vector magnitudes are always positive and so the second derivative of our objective function is always positive and our polynomial always has a minimum.

What about the case where Σ is not the identity matrix. Well because Σ is positive semi definite we can write it as:

$$\Sigma = \sqrt{\Sigma}\sqrt{\Sigma} \tag{14}$$

(see [6]) and so we can create the vectors:

$$\underline{x} = [\sqrt{\Sigma}]^{-1}\underline{x}' \qquad \underline{y} = [\sqrt{\Sigma}]^{-1}\underline{y}' \tag{15}$$

It is straightforward to show that by substituting (15) into (10) that with respect these new vectors, the endpoints of a new line segment, effectively, has the identity matrix covariance. Yet, the contrast optimisation is the same and we will find the same ρ as before. But, of course, by construction, because the covariance with respect to $\underline{x}$ and $\underline{y}$ is the identity the underlying polynomial has a unique minimum (we already demonstrated this). This argument suffices to show that (9) always has a unique minimum and so the best maximum contrast, brightness preserving, greyscale will lie on the endpoints of the line segment.

Thus we can dispense with QP. In our method we simply intersect two planes and bound the resulting line to lie within the unit cube. We check each endpoint in turn and choose the one that has maximum variance. It is worth stating, though it is self evident given the proof sketch above, that we get exactly the same result as QP but at negligible cost (intersecting planes is a simple non search algebraic operation).

As a final caveat we assume that there are two well defined planes to intersect. But, there are degenerate boundary cases e.g. if the mean in each of the R, G and B channels is the same. Then, in this cases the two planes are the same and their intersection is also a plane. So, in this case we would resort back to the QP solution. But, for the 100's of real images we possess, we have not found a degenerate case yet. But, this is a real possibility the reader should be aware of.

Experimental Results

We compared the maximum contrast outputs of our algorithm to results published by Grundland and Dodgson [8] and Alsam and Kolås's [1], see Figure 2. We implemented our own version of Alsam and Kolås's algorithm only so far as the greyscale conversion.

We compared the run-times of our closed form solution (implemented in MATLAB) and that of Qiu et al.'s QP method [10] (using MATLAB's own QP algorithm) to produce the solution coefficients. We ran both methods 6 times and took the mean time taken of the last 5 runs. Our closed form solution was seen to process 35 times faster.

Conclusion and Future Work

This paper presents a closed form solution to converting a colour image into a greyscale such that its contrast is maximised under brightness preserving constraints. To this end, we present a method using only simple analytical calculations, prior to which only a complex quadratic programming (QP) solution existed [10].

Our approach exploits the weighted sum method of colour to greyscale (Eq. (1)). The energy preserving (Eq. (3)) and brightness preserving (Eq. (4)) constraints are visualised using planar geometry whose intersection provides the line-segment solution of channel weights. Due to the quadratic form of the variance (Eq. (7)) and the posi semi definite nature of the covariance matrix (Eq. (5)), the solution for maximum greyscale image contrast will *always* exist at one of the end-points of the line segment.

Future work will consist of testing our reproduction for preference and also for 'information content'. This will enable us to provide statistical data on the effectiveness of our algorithm. Our psychovisual experiments will be reported in the final paper.

References

[1] Ali Alsam and Øyvind Kolås. Grey colour sharpening. In *Proceedings of 14th Color Imaging Conference*, pages 263–267. IS&T and SID, November 2006.

[2] Matthew Anderson, Ricardo Motta, Srinivasan Chandrasekar, and Michael Stokes. Proposal for a standard default color space for the internet: srgb. In *4th Color Imaging Conference*, pages 238–245. IS&T and SID, November 1996.

[3] Raja Bala and Reiner Eschbach. Spatial color-to-grayscale transform preserving chrominance edge information. In *Proceedings of 12th Color Imaging Conference*, pages 82–86. IS&T and SID, November 2004.

[4] David Connah, Graham D. Finlayson, and Marina Bloj. Seeing beyond luminance: a psychophysical comparison of techniques for converting colour images to greyscale. In *Proceedings of 15th Color Imaging Conference*, pages 336–341. IS&T and SID, 2007.

[5] Silvano Di Zenzo. A note on the gradient of a multi-image. *Computer Vision, Graphics, and Image Processing*, 33(1):116–125, January 1986.

[6] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*, page 149. 1996.

[7] Amy A. Gooch, Sven C. Olsen, Jack Tumblin, and Bruce Gooch. Color2gray: Salience-preserving color removal. *ACM Transactions on Graphics – Proceedings of ACM SIGGRAPH*, 24(3):634–639, July 2005.

[8] Mark Grundland and Neil A. Dodgson. The decolorize algorithm for contrast enhancing, color to grayscale conversion. Technical Report 649, University of Cambridge Computer Laboratory, October 2005.

[9] R. Hartley. *Linear and Nonlinear Programming*, chapter 12, page 183. Ellis Horwood LTD, 1985.

[10] Min Qiu, Graham D. Finlayson, and Guoping Qiu. Contrast maximizing and brightness preserving color to greyscale image conversion. In *Computer Graphics, Imaging and Visualization*, pages 347–351, June 2008.

[11] Karl Rasche, Robert Geist, and James Westall. Re-coloring images for gamuts of lower dimension. *Computer Graphics Forum*, 24(3):423–432, 2005.

[12] Diego A. Socolinsky and Lawrence B. Wolff. Multispectral image visualization through first-order fusion. *IEEE Transactions on Image Processing*, 11(8):923–931, August 2002.

Author Biography

Graham D. Finlayson obtained his BSc in Computer Science from the University of Strathclyde (Glasgow, Scotland) in 1989. He then pursued his graduate education at Simon Fraser University (Vancouver, Canada) where he was awarded his MSc and PhD degrees in 1992 and 1995 respectively. From August 1995 until September 1997, Dr. Finlayson was a Lecturer in Computer Science at the University of York (York, UK) and from October 1997 until August 1999 he was a Reader in Colour Imaging at the Colour and Imaging institute, University of Derby (Derby, UK). In September 1999 he was appointed a Professor in the School of Computing Sciences, University of East Anglia (Norwich, UK).

Toby N. Matheson received his BSc in Physics in 2003 and his MSc in applied optics in 2005 from the University of Reading (UK). From 2005-2011 he conducted materials research at the University of Southampton and the University of Surrey (UK) before discovering his passion for image processing. Since 2011 he has been a PhD student at the University of East Anglia (Norwich, UK) under the supervision of Prof. Graham D. Finlayson.

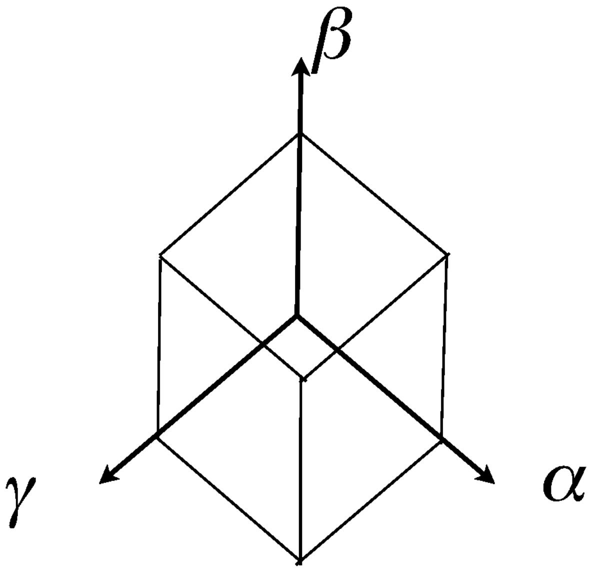

(a) Visualising (2)

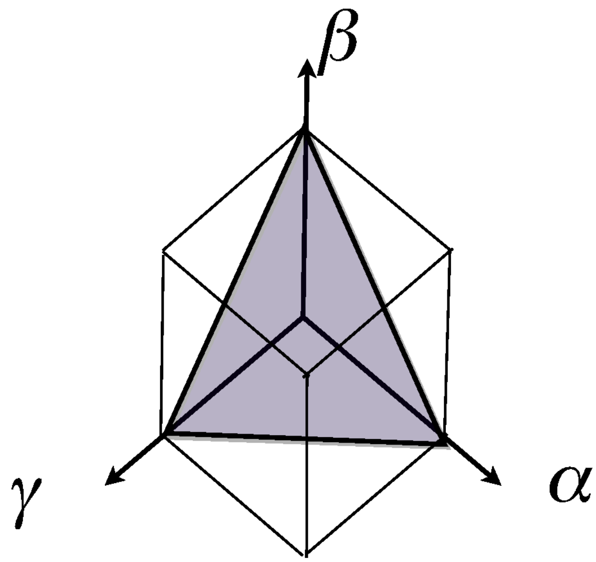

(b) Visualising (2) and (3)

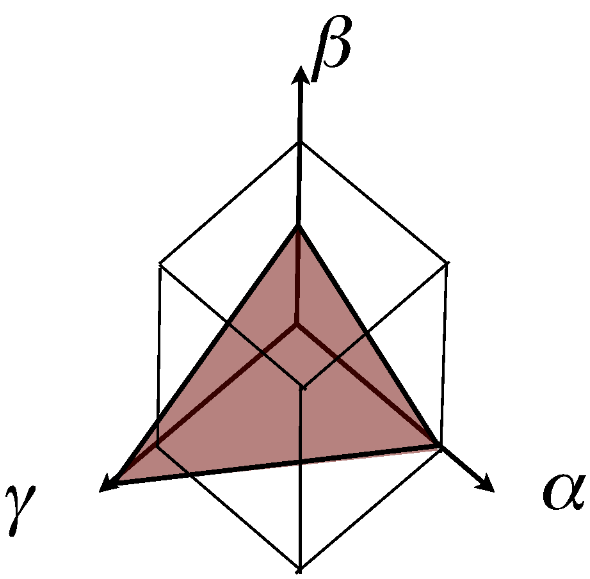

(c) Visualising (2) and (4)

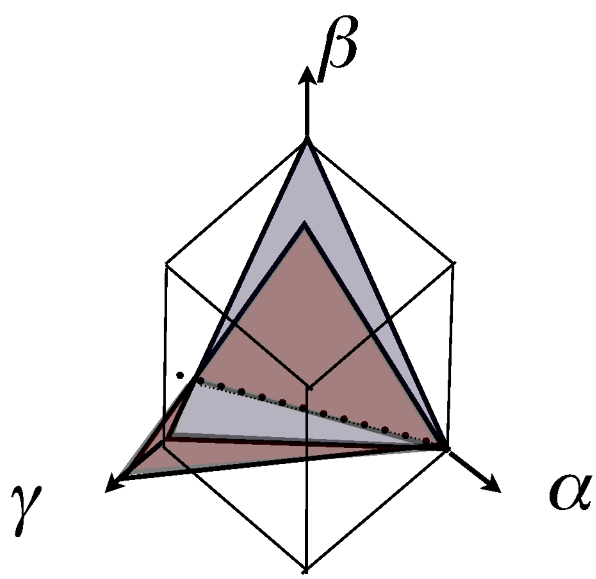

(d) Visualising (2), (3) and (4)

Figure 1: Visual representation of the closed form solution

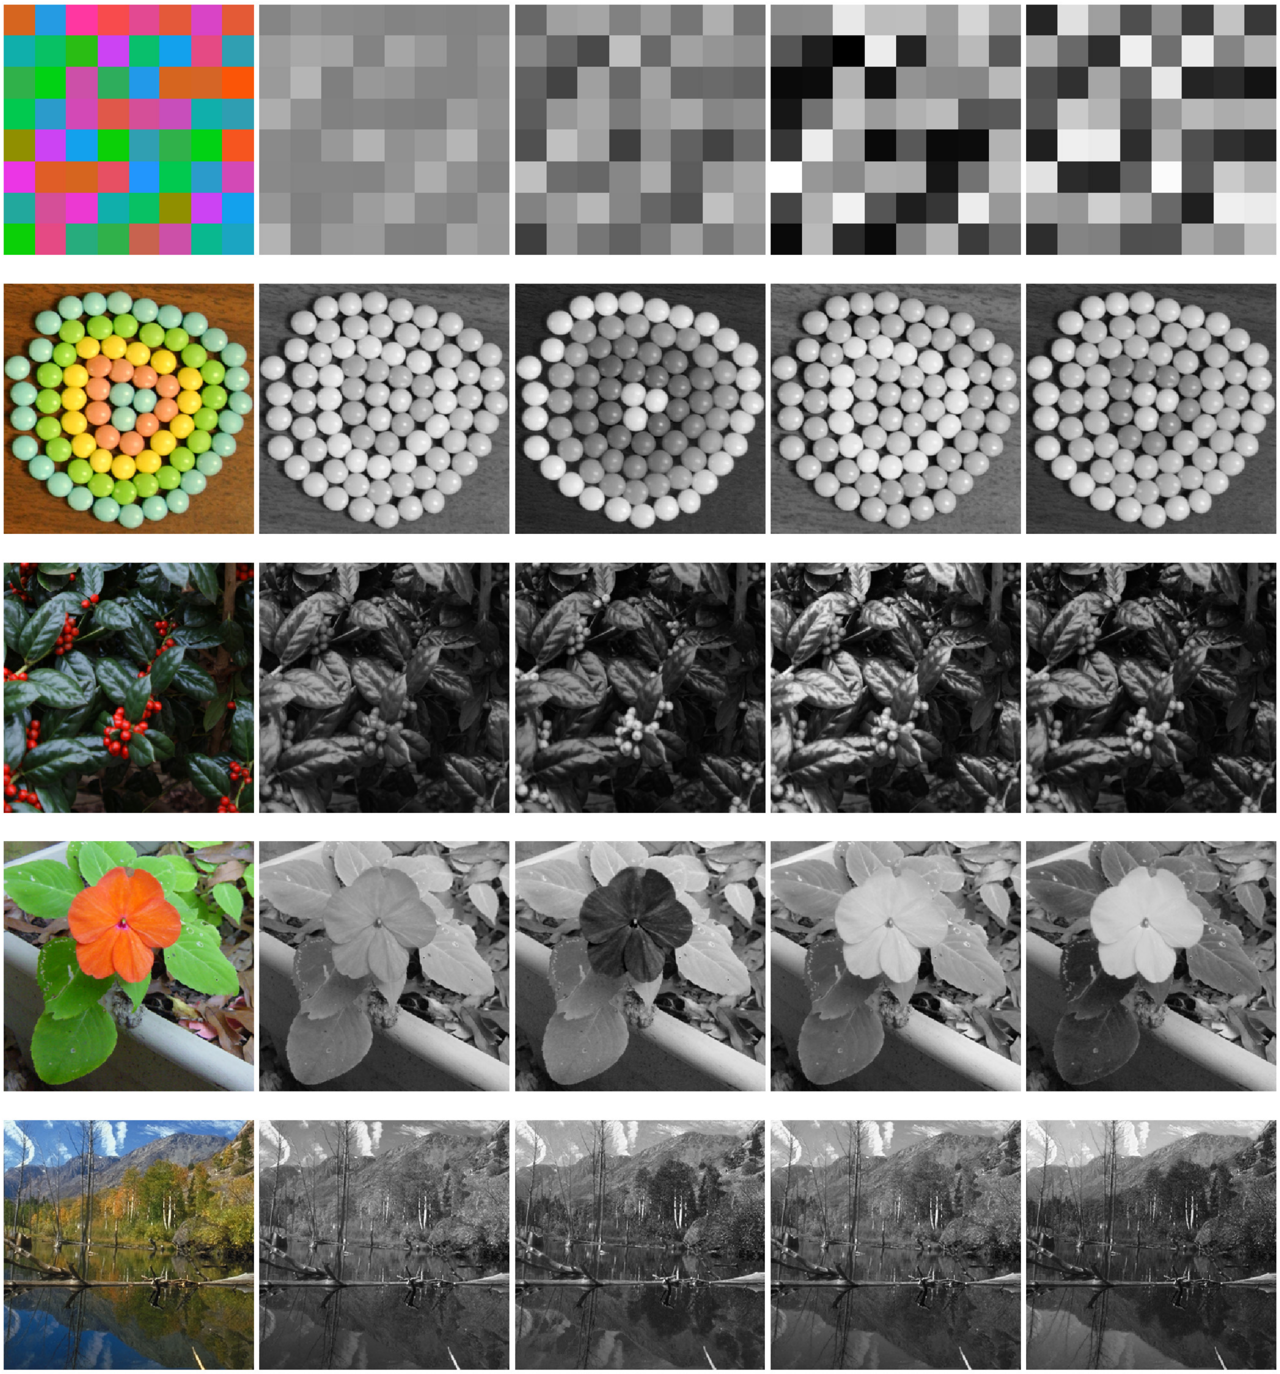

RGB | Luminance | Grundland and Dodgson | Alsam and Kolås | Our Method

Figure 2: Comparison of colour to greyscale conversion between the luminance channel in sRGB colour space, Grundland and Dodgson, Alsam and Kolås, and our own